

# Distributed DTX Alignment with Memory

Hauke Holtkamp, Guido Dietl

DOCOMO Euro-Labs  
D-80687 Munich, Germany  
Email: {holtkamp, auer}@docomolab-euro.com

Harald Haas

Institute for Digital Communications  
Joint Research Institute for Signal and Image Processing  
The University of Edinburgh, EH9 3JL, Edinburgh, UK  
E-mail: h.haas@ed.ac.uk

November 7, 2014

## Abstract

This paper addresses the assignment of transmission and sleep time slots between interfering transmitters with the objective of minimal power consumption. In particular, we address the constructive alignment of Discontinuous Transmission (DTX) time slots under link rate constraints. Due to the complexity of the combinatorial optimization problem at hand, we resort to heuristic assignment strategies. We derive four time slot alignment solutions (sequential alignment, random alignment, p-persistent ranking and DTX alignment with memory) and identify trade-offs. One solution, DTX alignment with memory, addresses trade-offs of the other three by maintaining memory of past alignment and channel quality to buffer short term changes in channel quality. All strategies are found to exhibit similar convergence behavior, but different power consumption and retransmission probabilities. DTX alignment with memory is shown to achieve up to 40% savings in power consumption and more than 20% lower retransmission probability than the State-Of-The-Art (SotA).

## 1 Introduction

To reduce their power consumption, future Base Stations (BSs) will use very short sleep modes—called Discontinuous Transmission (DTX)—which

interrupt their transmission [?]. DTX can be employed for energy saving when a communication system has spare capacity, *e.g.* at night in cellular networks. It can be easily integrated into Orthogonal Frequency Division Multiple Access (OFDMA) systems without incurring delays by not scheduling resources for transmission during one or more time slots and then setting the transmitter to DTX mode instead for the duration of these unscheduled time slots. For the transmitter, this is a local power saving measure. When seen from a network perspective, these regular interruptions in transmission can be aligned to reduce interference and power consumption. Such alignment has not been studied in detail, but is related to the field of channel allocation.

Generally, in channel allocation, the challenge is to assign frequency channels to different cells or links of a network such that mutual interference is minimized. To be applicable for future small-cell cellular networks, such channel allocation needs to be distributed, uncoordinated and dynamic. Distributed operation prevents the delay and backhaul requirements imposed by a central controller. If the allocation is also uncoordinated, it does not require message exchange through a backhaul of limited capacity. Also, as networks, channels and traffic change rapidly, the allocation must be highly dynamic and update often. Problems of this type are known to be NP-hard [?]. Therefore, different suboptimal approaches have been proposed. The simplest approach to distributing channels over a network is a fixed assignment in which it is predefined which links will use which channels, *e.g.* see [?]. However, this is inflexible to changing or asymmetric traffic loads. For flexibility, dynamic channel allocation methods are required. A simple dynamic channel allocation method, Sequential Channel Search [?], assigns channels with sufficient quality in a predefined order. This technique allows adjusting the number of channels flexibly, but is sub-optimal due to its strong channel overlap between neighbors. Taking the channel quality into consideration by measurement is proposed in the Minimum Signal-to-Interference-and-Noise-Ratio (SINR) method [?]. This technique offers increased spectral efficiency, but can lead to instabilities when allocations happen synchronously. Ellenbeck *et al.* [?] introduce methods of game theory and respond to the instabilities by adding p-persistence to avoid simultaneous bad player decisions. However, the proposed method only applies to single user systems and does not address target rates. Dynamic Channel Segregation [?] introduces the notion of memory of channel availability, allowing a transmitter to track which channels tend to be favorable for transmission. However, the algorithm only applies to sequential channel decisions based on an idle or busy state in circuit switched networks

and cannot be applied to the concurrent alignment of channels as is required in OFDMA networks.

In this paper, we combine the findings from past channel allocation research and adapt them for aligning DTX time slots in an OFDMA frame. We derive four distributed uncoordinated dynamic time slot alignment strategies for the cellular downlink, in which BSs independently prioritize time slots for transmission while maintaining system stability. The remainder of the paper is structured as follows. Section 2 formulates the system model and the problem at hand. The considered four alternative solutions are described in Section 3. Findings obtained from simulation are presented in Section 4. The paper is concluded in Section 5.

## 2 System model and problem formulation

We make the following assumptions about the network. BSs in a reuse-one OFDMA cellular network schedule a target rate,  $B_k$ , per OFDMA frame to be transmitted to each mobile  $k = \{1, \dots, K\}$ . All time slots are available for scheduling in all cells. Mobiles report perceived SINR from the previous OFDMA frame,  $s_{n,t,k}$ , on subcarrier  $n = \{1, \dots, N\}$ , time slot  $t = \{1, \dots, T\}$  to their associated BS. BSs have a DTX mode available for each time slot during which transmission and reception are disabled and power consumption is significantly reduced to  $P_S$  compared to power consumption during transmission, which is a function of the power allocated to each Resource Block (RB),  $\rho_{Tx}$ . DTX is available fast enough to enable it within individual time slots of an OFDMA transmission such that transmission time slots are not required to be consecutive. A BS can schedule no transmission in one or more time slots and go to DTX mode instead. Scheduling a DTX time slot in one BS reduces the interference on that time slot to all other BSs. The OFDMA frames of all BSs are assumed to be aligned such that the interference over one time slot and subcarrier is flat and that all BSs can perform alignment operations in synchrony. The OFDMA frame consists of  $NT$  RBs. The channel is subject to block fading. Interference is treated as noise.

To describe the resource allocation problem formally, we define the function

$$\begin{aligned} \Pi : \{1, \dots, T\} \times \{1, \dots, N\} &\rightarrow \{0, 1, \dots, K\} \\ (n, t) &\mapsto k, \end{aligned} \tag{1}$$

which maps to each RB a user  $k$  or 0, where 0 indicates that the RB is not scheduled.

Each cell tries to minimize its power consumption by scheduling time slots for transmission and sleep. With  $r_{n,t,k}$  the capacity of resource  $(n, t)$  if it is scheduled to  $k$ , the optimization problem for the total BS power consumption,  $P_{\text{total}}$ , is as follows.

$$\underset{\Pi}{\text{minimize}} \quad P_{\text{total}} = P_{\text{S}} \frac{T_{\text{S}}}{T} + \rho_{\text{Tx}} N_{\text{Tx}} + P_0 \left( \frac{T - T_{\text{S}}}{T} \right) \quad (2a)$$

subject to

$$N_{\text{Tx}} = \quad (2b)$$

$$|\{n \in \{1, \dots, N\} | (\Pi(n, t) \neq 0 \quad \forall t \in \{1, \dots, T\})\}|$$

$$T_{\text{S}} = \quad (2c)$$

$$|\{t \in \{1, \dots, T\} | (\Pi(n, t) = 0 \quad \forall n \in \{1, \dots, N\})\}|$$

$$B_k \leq \sum_{\{(n,t) | \Pi(n,t)=k\}} r_{n,t,k} \quad \forall k. \quad (2d)$$

With  $|\cdot|$  the cardinality operator,  $T_{\text{S}}$  is the number of time slots in which for all  $n$ , no user is allocated and the BS can enable DTX and  $N_{\text{Tx}}$  is the number of RBs that are scheduled for transmission. The constraint (2d) provides the rate guarantee for each user.

The difficulty lies in finding the mapping  $\Pi$ . Under a brute force approach, there exist  $(K + 1)^{NT}$  possible combinations. Due to the large number RBs present in typical OFDMA systems like Long Term Evolution (LTE), the computation of the solution is infeasible. Consequently, we resort to heuristic methods in the next section.

Some of the methods compared in this paper make use of a time slot ranking. In other works, *e.g.* [?], ranking is proposed to be based on SINR. However, in a multi-user OFDMA system each subcarrier and mobile terminal have a different SINR, thus generating a problem of comparability between time slots. Consequently, we propose to compare time slots by their hypothetical sum capacity,  $B_t$ , over all mobiles and subcarriers with

$$B_t = \sum_{k=1}^K \sum_{n=1}^N \log_2 (1 + s_{n,t,k}). \quad (3)$$

### 3 DTX alignment strategies

In this section, we derive four strategies to tackle the problem at hand which differ in performance and complexity.

### 3.1 Sequential alignment

This strategy is to always allocate as many time slots for transmission as required in sequential order and set the remainder to DTX. This leads to strongest overlap and consequently highest interference on the first time slot and lowest overlap and possibly no transmissions at all on the last time slot. This strategy does not make use of the available channel quality information and provides valuable insights into questions of stability, reliability and convergence. It serves as a deterministic upper bound.

### 3.2 Random alignment

Random alignment refers to randomly selecting transmission time slots for every OFDMA frame and setting the remainder to DTX. This strategy provides a reference for achievable gains, allows to assess the worst case effects of randomness and represents the state of the art in today's unsynchronized unaligned networks.

### 3.3 P-persistent ranking

The synchronous alignment of uncoordinated BSs can lead to instabilities, when neighboring BSs—perceiving similar interference information—schedule the same time slots for transmission. This leads to oscillating scheduling which never reaches the desired system state [?]. To address this problem, we introduce p-persistence to break the unwanted synchrony by only changing established DTX schedules with a probability  $p = 0.3$ . In initial tests, the value of  $p$  did not have a significant effect on the power consumption. The exact tuning of  $p$  is out of the scope this paper. P-persistent ranking first ranks time slots by their sum capacity as in (3), schedules time slots in that order, but only applies this new selection with probability  $p$ . Otherwise, the schedule from the last iteration remains active. After ranking, time slots are selected for transmission in order of  $B_t$ . The remainder of time slots is set to DTX.

### 3.4 Distributed DTX alignment with memory

To counter oscillation and achieve a convergent network state we introduce memory of past schedules into the alignment process. The rationale behind the distributed DTX alignment with memory algorithm is as follows. Taking into account current time slot capacities,  $B_t$ , past scores and slot allocations, each BS first updates the internal score of each time slot and then returns the

priority of time slots by score. In case of equal scores, time slots are further sub-sorted by  $B_t$ . The score is updated in integers. All time slots which were used for transmission in the previous OFDMA frame receive a score increment of one. Furthermore, the time slot with highest  $B_t$ ,  $R_0$ , receives an increment of one. Time slots which were not used for transmission in the previous OFDMA frame receive a decrement of one, except for  $R_0$ . Scores have an upper limit  $\psi_{ul}$  beyond which there is no increment and a lower limit  $\psi_{ll}$  below which there is no decrement. The difference between  $\psi_{ul}$  and  $\psi_{ll}$  can be interpreted as the depth of the memory buffer.

---

**Algorithm 1** Distributed DTX alignment with memory

---

**Input:**  $\psi, \Upsilon_u$

```

1:  $R \leftarrow \text{sort-desc-by-capacity}(\Upsilon)$ 
2: for all  $v$  in  $\Upsilon_u$  do
3:   if  $\psi(v) < \psi_{ul}$  then
4:      $\psi(v) \leftarrow \psi(v) + 1$ 
5:   end if
6: end for
7: for all  $v$  in  $\Upsilon_{uu} \setminus \{R_0\}$  do
8:   if  $\psi(v) > \psi_{ll}$  then
9:      $\psi(v) \leftarrow \psi(v) - 1$ 
10:  end if
11: end for
12:  $\psi(R_0) \leftarrow \psi(R_0) + 1$ 
13:  $V \leftarrow \text{sort-desc-by-score}(R, \psi, \Upsilon)$ 
14: return  $\psi, V, \Upsilon_u$ 
```

---

The described steps are summarized in Algorithm 1 with scoring map  $\psi$ , ranking tuple  $R$ , and priority tuple  $V$ . The set of used time slots  $\Upsilon_u$  and unused time slots  $\Upsilon_{uu}$  make up the set of time slots  $\Upsilon$ , the cardinality of which is  $T$ . The function 'sort-desc-by-capacity( $\Upsilon$ )' returns a list of time slots ordered descending by  $B_t$ . The function 'sort-desc-by-score( $R, \psi, \Upsilon$ )' returns a list of time slots ordered primarily by descending score and secondarily by descending  $B_t$ . After the execution of Algorithm 1, the first  $T_{Tx}$  time slots in  $V$  are scheduled for transmission.

In the following, we illustrate the iterations over three OFDMA frames of Algorithm 1 for a system with three time slots. In the example, the algorithm delays the change of  $\Upsilon_u$  from  $c$  to  $b$  to buffer scheduling changes. The iteration begins with arbitrarily chosen  $\Upsilon = \{a, b, c\}$ ,  $\psi = \{a : 0, b : 2, c : 5\}$ :

1.  $\Upsilon_u = \{c\}$ ,  $R = (b, c, a)$   
 $\rightarrow \psi = \{a : 0, b : 3, c : 5\}$ ,  $V = (c, b, a)$
2.  $\Upsilon_u = \{b, c\}$ ,  $R = (b, c, a)$   
 $\rightarrow \psi = \{a : 0, b : 5, c : 5\}$ ,  $V = (b, c, a)$
3.  $\Upsilon_u = \{b\}$ ,  $R = (b, a, c)$   
 $\rightarrow \psi = \{a : 0, b : 5, c : 4\}$ ,  $V = (b, c, a)$

When applied, Algorithm 1 strongly benefits time slots which were used for transmission in the past ( $b, c$  in step 1 of the example). These time slots tend to repeatedly receive score increments until they all have maximum score  $\psi_{ul}$  ( $b, c$  in step 2 of the example). When the score is equal for some time slots ( $b, c$  in step 2 of the example), the ranking is based on  $B_t$ . A time slot which was used for transmission and has highest  $B_t$ ,  $R_0$ , receives the highest increment (slot  $b$  in step 2 of the example). When a time slot has highest  $B_t$ , but was not selected for transmission in the previous OFDMA frame, it receives a score increment, but is not guaranteed to be used for transmission (slot  $b$  in step 1 of the example). When a time slot repeatedly has highest  $B_t$ , it reaches  $\psi_{ul}$  ( $b$  in the example). The algorithm thus buffers short term changes in the channel quality setting in favor of long term time slot selection.

## 4 Results

We analyze the four alignment strategies using computer simulations with regard to power consumption, convergence, reliability of delivered rates and algorithmic complexity after the introduction of the simulation environment and the RB scheduling scheme.

### 4.1 Simulation environment and RB scheduling

The four strategies were tested in a network simulation with 19-cell hexagonal arrangement with uniformly distributed mobiles and fixed target rates per mobile. Data was collected only from the center cell which is thus surrounded by two tiers of interfering cells. Power consumption of a cell is modeled as a function of transmission power as described in [?]. Table 4.1 lists additional parameters used which approximate an LTE system. The simulation is started with the assumption of full transmission power on all resources with a power consumption of 350 W per cell as a worst-case initial configuration.

Parameter	Value
Carrier frequency	2 GHz
Intersite distance	500 m
Pathloss model	3GPP UMa, NLOS, shadowing [?]
Shadowing standard deviation	8 dB
Bandwidth	10 MHz
Transmission power per RB	0.8 W
Thermal noise temperature	290 K
Interference tiers	2 (19 cells)
Mobile target rate	2 Mbps and 3 Mbps
OFDMA subframes (time slots)	10
Subcarriers	50
Mobiles	10
$\psi_{ul}$	5
$\psi_{ll}$	0
Power model factors [?] (idle; load factor; DTX)	200 W; 3.75; 90 W

Table 1: Simulation Parameters

In order to assess the performance of the four strategies, it is necessary to make assumptions about how the individual RBs are scheduled within a time slot. In order to avoid masking effects of the algorithms under test, we have applied a sequential RB allocation. Sequential resource allocation is performed after the DTX alignment step is completed and allocates as many bits to an OFDMA resource block as possible according to the Shannon capacity, followed in order by the next resource in the same time slot (sequentially), until the target rate has been scheduled to each mobile. Time slots are scheduled in the order provided by each of the four strategies. This sequential resource scheduler deliberately omits the benefits of multi-user diversity. This leads to underestimating achievable rates in simulation compared to a system which exploits multi-user diversity, but allows a fair comparison of the quality of different DTX alignment algorithms.

## 4.2 Power consumption

To assess achievable power savings and the dynamic adaptivity over a large range of cell loads, we inspect the cell total power consumption in Fig. 1. At low load very few RBs are required to deliver the target rate and more time slots can be scheduled for DTX than at high traffic loads, leading to






Figure 1: BS power consumption over different cell sum rates.

monotonously rising power consumption over increasing target rates for all alignment methods.

Sequential alignment causes the highest power consumption over any target rate with an almost linear relationship between user target rates and power consumption. Sequential alignment consumes high power, as it schedules many RBs to achieve the target rate due to the high interference level present.

This power consumption is significantly lower for the State-Of-The-Art (SotA), random alignment. The randomness of time slot alignment creates a much lower average interference than with sequential alignment allowing more data to be transmitted in each RB. As fewer RBs are required, less power is consumed.

P-persistent ranking and DTX with memory both achieve similarly low power consumption of up to 40% less than random alignment. The relationship between power consumption and target rate is noticeably non-linear, as it is flat at low target rates and grows more steeply at high target rates. This behavior is caused by the low interference level these strategies manage to create. Only at high rates, when the number of sleep time slots has to be significantly decreased, does the interference increase, leading to higher power consumption.

Also noteworthy is the fact that at 1 Mbps and 3 Mbps, random alignment performs nearly as good as p-persistent ranking. At these extreme points the network is almost unloaded and almost fully loaded, respectively. Consequently, either most time slots are scheduled for DTX or none, leaving very little room for optimization compared to randomness. The largest potential for time slot alignment for power saving is in networks which are medium loaded. Under medium load, the number of transmission and DTX time slots is similar, causing the effects of alignment to be most pronounced.




Figure 2: BS power consumption over OFDMA frames at 1 Mbps per mobile.




Figure 3: BS power consumption over OFDMA frames at 2 Mbps per mobile.

### 4.3 Convergence

Another relevant aspect is the convergence of the network to a stable state. As each BS makes iterative adjustments to its selection of time slots for transmission, the speed of convergence as well as the convergence to a stable point of operation are relevant criteria. The effect of the iterative execution of the four strategies on the cell power consumption is illustrated in Fig. 2. Power consumption is found to converge to a stable value within within six OFDMA frames (alignment iterations). All strategies converge to the average power consumption values shown in Fig. 1. The simulation starts from a worst-case schedule of transmission on all RBs and then iteratively schedules time slots for transmission and DTX. In the case of 1 Mbps per user, one transmission time slot is sufficient to schedule the target rate. With transmissions only taking place during one time slot, there is very little difference between a random alignment and p-persistent ranking or DTX with memory. At 2 Mbps per user, see Fig. 3, random alignment occasionally causes higher interference than p-persistent ranking or DTX with memory, leading to higher power consumption. Also, p-persistent ranking converges more slowly than DTX alignment with memory.




Figure 4: Retransmission probability over targeted rate.

#### 4.4 Reliability

An important aspect in dynamical systems with target rates is that scheduled target rates cannot always be fulfilled. As RB scheduling is based on channel quality information which was collected in the previous OFDMA frame, the actual channel quality during transmission may differ, leading to lower than expected rates. Thus, although certain target rates are scheduled and although the system is not fully loaded, a BS may fail to deliver the targeted rate and require retransmission of some RBs. We assess this metric by considering the retransmission probability for each strategy. The results are shown in Fig. 4 against a range of user target rates.

Easiest to interpret is sequential alignment which does not require retransmission for target rates up to 2.3 Mbps due to its determinism. The increase in the retransmission probability at high rates is not caused by a failure of the alignment, but by system overload. When high rates are combined with bad channel conditions, the system may be unable to deliver the target data rate, independent of the alignment strategy. This increase of the retransmission probability at high rates is present for all alignment strategies and constitutes outage.

The retransmission probability is highest for random alignment. This is caused by the strong fluctuation of interference under randomized scheduling. Channel quality measurements used for RB scheduling are of very little reliability, as the interference changes quickly due to random time slot alignment, resulting in increased retransmission probability.

P-persistent ranking performs slightly better than random alignment at 1 Mbps target rate. At 2 Mbps per user, where the alignment potential is highest, oscillation causes the highest retransmission probability.

DTX alignment with memory achieves a much lower retransmission probability in the range of 15% to 20%, due to the reduction in interference fluctuation introduced by the memory score.

## 4.5 Complexity

With regard to complexity, sequential and random alignment are of minimal complexity. These strategies involve no algorithmic decision-making on the set of transmission time slots. P-persistent ranking requires one ranking and time slot selection with probability  $p$  per iteration. The highest complexity is present in DTX with memory, which requires two executions of the time slot sort, one for the generation of the score and one for the output of the ranking. Although DTX with memory comprises the highest complexity in comparison, these operations pose a small burden on modern hardware as the number of time slots is typically small. For example, typical LTE systems are designed with 10 subframes.

## 4.6 Interpretation

To conclude our evaluation, we have found that DTX with memory provides the best results. Under the present assumptions, it provides both a lower power consumption and lower retransmission probability than the SotA and p-persistent ranking. Future DTX capable networks can and should exploit this alignment potential.

## 5 Conclusion

In this paper, we have studied the constructive uncoordinated and distributed alignment of DTX time slots between interfering OFDMA BSs for power-saving under rate constraints. We first established the open problem formally. Due to its complexity, we have approached it with four alternative heuristic strategies. One of these strategies, DTX with memory, is an original contribution of this paper and introduces memory to overcome short-term fluctuations and networks oscillations. The performance of the four strategies was tested in simulation. It was found that power consumption can be reduced significantly, especially at medium cell loads. All strategies converge within six OFDMA frames. DTX with memory reduces power consumption up to 40% compared to the SotA, combined with around 20% reduction in retransmission probability.